

A Bi-dimensional Regression Tree Approach to the Modeling of Gene Expression Regulation

Jianhua Ruan¹ and Weixiong Zhang^{1,2*}

¹Department of Computer Science and Engineering, and ²Department of Genetics, Washington University in St Louis, St Louis, MO 63130, USA

ABSTRACT

Motivation: The transcriptional regulation of a gene depends on the binding of cis-regulatory elements on its promoter to some transcription factors and the expression levels of the transcription factors. Most existing approaches to studying transcriptional regulation model these dependencies separately, i.e., either from promoters to gene expression or from the expression levels of transcription factors to the expression levels of genes. Little effort has been devoted to a single model for integrating both dependencies.

Results: We propose a novel method to model gene expression using both promoter sequences and the expression levels of putative regulators. The proposed method, called bi-dimensional regression tree (BDTree), extends a multivariate regression tree approach by applying it simultaneously to both genes and conditions of an expression matrix. The method produces hypotheses about the condition-specific binding motifs and regulators for each gene. As a side-product, the method also partitions the expression matrix into small submatrices in a way similar to bi-clustering. We propose and compare several splitting functions for building the tree. When applied to two microarray data sets of the yeast *Saccharomyces cerevisiae*, BDTree successfully identifies most motifs and regulators that are known to regulate the biological processes underlying the data sets. Comparing to an existing algorithm, BDTree provides a higher prediction accuracy in cross-validations.

Availability: The software is available upon request from the authors.

Contact: jruan@cse.wustl.edu, zhang@cse.wustl.edu

Supplementary Information: <http://cic.cs.wustl.edu/bdtree/>

1 INTRODUCTION

The complex function of a living cell is controlled by regulating the expression of specific genes at several levels. One of the most important and best understood regulation mechanisms is at the transcriptional level, where the expression of a gene is mediated by the binding of transcription factors (TFs) to specific DNA sequences in the promoter region of the gene. Two basic assumptions are often made when studying transcriptional regulation: first, the expression of a gene is determined by the binding sites of specific TFs on its promoter; second, the expression of a gene is a function of the concentration of specific TFs around its promoter. Based on the two assumptions, two distinct classes of approaches have been proposed in studying transcriptional regulation.

The first class of approaches attempted to build quantitative or qualitative models to associate gene expression levels with putative binding motifs on their promoter sequences (Fig. 1 boxes A and




Fig. 1. Relationships between our method and previous methods. The bottom right matrix represents gene expression levels. The bottom left matrix represents motif occurrences on promoter sequences. The top matrix is the expression levels of regulators. Box A, the expression levels of multiple genes under a single condition are modeled by the motifs on their promoters (e.g., Bussemaker *et al.*, 2001). Box B, the expression levels of multiple genes under multiple conditions are modeled by the motifs on their promoters (Phuong *et al.*, 2004). Box C, the expression levels of a single gene under multiple conditions are modeled by the expression levels of putative regulators (Soinov *et al.*, 2003). Box D, the expression levels of multiple genes under multiple conditions are modeled by the expression levels of putative regulators (Segal *et al.*, 2003). Box E, the expression levels of multiple genes under multiple conditions are modeled by the motifs on their promoters and the expression levels of putative regulators (Middendorf *et al.*, 2004 and our method).

B). Several approaches of this type have been proposed within the classification and regression framework. Bussemaker *et al.* (2001) and others (Keles *et al.*, 2002; Conlon *et al.*, 2003) modeled the expression levels of genes as a linear regression of putative binding motifs, and applied feature selection techniques to find the most significant motifs. Hu *et al.* (2000) and ourselves (Ruan and Zhang, 2004) used decision trees to find motif combinations that best separate two sets of genes. Beer and Tavazoie (2004) built probabilistic graphical models, e.g., Bayesian networks, to explain gene expression patterns from motifs. Phuong *et al.* (2004) applied multivariate regression trees to model the transcriptional regulation of gene expression over several time points simultaneously.

The second class of approaches have been proposed to model gene expression levels from the expression levels of other genes, i.e., TFs and other regulators (Fig. 1 boxes C and D). For example, Soinov *et al.* (2003) used decision tree to identify possible regulators for several cell-cycle genes individually. Segal *et al.* (2003) proposed a more sophisticated procedure suitable for whole-genome analysis. The method first clusters genes according to their expression patterns, and then builds a regression tree for each cluster to represent

*to whom correspondence should be addressed

their common regulation program. The procedure then iteratively refines the clusters and the trees.

Middendorf *et al.* (2004) recently introduced a method that combines the previous two classes of approaches. Their method models gene expression levels from both putative binding motifs on promoter sequences and the expression levels of putative regulators (Fig. 1 box E). Here we propose a method that also falls into this category. Although our method has the same schematic representation as their method, the underlying modeling rationales are very different, which we will compare in detail in Section 4.

Our method, called bi-dimensional regression tree or BDTree for short, is an extension to the multivariate regression tree approach of Segal (1992) and Phuong *et al.* (2004). Breiman *et al.* (1984) first introduced the univariate regression tree approach to recursively partition instances into groups, where the instances in each group have similar attribute values and responses. Segal (1992) extended the method to handle multiple responses, so that the instances in each group have a similar pattern of responses across multiple conditions. The basic idea of our method, as suggested by its name, is to extend the multivariate regression tree approach to both dimensions of the expression matrix (see Fig. 1). On one dimension, each gene is treated as an instance, where the attributes are the binding motifs on its promoter sequence, and the responses are its expression levels across the conditions. Genes are partitioned so that those in the same subset have common binding motifs and similar expression patterns across the conditions. On the other dimension, each condition is treated as an instance, where the attributes are the expression levels of candidate regulators under the condition, and the responses are the expression levels of genes under that condition. Conditions are partitioned so that the expression levels of a gene under each subset of conditions are similar.

The way of partitioning genes and conditions in BDTree is analogous to bi-clustering (Cheng and Church, 2000). However, the partitioning in BDTree is supervised by some intrinsic attributes of the genes and conditions, i.e., the binding motifs on gene promoters and the expression levels of regulators under the conditions. In contrast, in bi-clustering, the partitioning is unconstrained by those attributes. As a result, the model learned by BDTree is both exploratory and predictive. It suggests a set of testable hypotheses of condition-specific binding motifs and regulators for the genes in each cluster, and can also be used to predict the expression levels of unseen genes under unseen conditions, given appropriate attributes of the genes and conditions.

The rest of the paper is organized as follows. The next section first introduces the univariate regression tree and its multivariate extension, and then describes the bi-dimensional multivariate regression tree approach. Section 3 presents some experimental results from applying the method to the yeast cell-cycle and stress response data. In the last section we discuss the differences and relationships between BDTree and several related methods.

2 ALGORITHM

2.1 Univariate regression trees

Here we give a brief overview of the univariate regression tree method and refer the reader to Breiman *et al.* (1984) for details. Suppose that there are p attributes X_1, X_2, \dots, X_p and a response Y . The values of the attributes and responses are observed for m instances: $\{(\mathbf{x}_i, y_i)\}$, where $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$, for $i = 1, \dots, m$. Here

the responses are real values. We restrict all attributes to be real values for convenience of the discussion, although the method can handle categorical or mixed values. In the context of transcriptional regulation, each gene is an instance, the attributes are motifs, and the response is the gene expression level under a single condition (Fig. 1, box A).

In the classic CART (Classification Analysis and Regression Tree) program of Breiman *et al.* (1984), a greedy search algorithm is used to construct a binary regression tree. The basic algorithm is as follows.

1. Initially there is only the root node containing all instances.
2. If the current node has not met the stopping criterion, examine every possible binary split of the instances within the node based on each attribute X_i , $i = 1, \dots, p$, such that the attribute values for all the instances in one subset are smaller than those in the other subset.
3. Choose the best split to maximize an objective function, and create two child nodes for the current node.
4. Repeat step 2 and 3 for each child node.

To build a regression tree, three rules need to be specified: a splitting rule that defines the best split, a stopping rule that determines when the splitting should terminate, and a third rule to prune certain branches of the tree after the tree is built. Here we only discuss the splitting rule, while the other two will be discussed after we introduce the BDTree method.

The goal of a split is to produce child nodes as homogeneous as possible with respect to the responses. A frequently used criterion is the least-square rule which aims at minimizing the sum-of-squares of responses within each node. Let r denote a node of the tree and n_r denote the number of instances in r . That is, r contains a subset of the indices $\{1, \dots, m\}$. The within-node sum-of-squares is given by

$$SS(r) = \sum_{i \in r} (y_i - \bar{y})^2, \quad (1)$$

where $\bar{y} = \sum_{i \in r} \frac{y_i}{n_r}$. The gain of a split that partitions r into two child nodes r_1 and r_2 is given by

$$G(r, r_1, r_2) = SS(r) - SS(r_1) - SS(r_2). \quad (2)$$

The best split is determined by an attribute X_i and a threshold T such that $G(r, r_1, r_2)$ is maximized and that the value of X_i for every instance in r_1 is less than T while that for every instance in r_2 is no less than T . To find the best split, all possible thresholds for each attribute are tested, and the split with the highest gain is chosen.

2.2 Multivariate regression trees

It is not uncommon to encounter domains where the responses are observed under multiple conditions, that is, the response of an instance is also a vector: $\mathbf{y}_i = (y_{i1}, \dots, y_{in})$, where n is the number of conditions and y_{ij} is the response of the i th instance under the j th condition. For example, in DNA microarrays, gene expression levels are typically recorded for several time points or experimental conditions.

A naive solution for this situation is to build a regression tree for each condition separately. However, it is difficult to combine

multiple trees. Segal (1992) introduced a multivariate regression tree method to construct a single tree to model multiple responses simultaneously. He generalized the within-node sum-of-squares in Equation (1) as follows:

$$SS1(r) = \sum_{i \in r} (\mathbf{y}_i - \bar{\mathbf{y}}) \mathbf{V}^{-1} (\mathbf{y}_i - \bar{\mathbf{y}})^t, \quad (3)$$

where \mathbf{y}_i is the vector of responses for the i th instance, \mathbf{V} is the model covariance matrix of \mathbf{y}_i , and $\bar{\mathbf{y}}$ is the average of \mathbf{y}_i within node r . With $SS(r)$ defined, the gain function remains the same as Equation (2) and the recursive algorithm proceeds to split the instances as in the case with a single response.

Phuong *et al.* (2004) applied the multivariate regression tree method to gene expression data by treating genes as instances, where the numbers of occurrences of motifs in promoters are the attribute values, and the expression levels at different conditions are the multivariate responses (Fig. 1, box B). As noted by Segal (1992) and Phuong *et al.* (2004), the multivariate regression tree method is intermediary between classification and clustering. The responses of different instances can be written as a matrix $\mathcal{Y} = (y_{ij})$, where i is the index of an instance and j is the index of a condition. The multivariate regression tree partitions the matrix into sub-matrices, where each sub-matrix contains all the columns (conditions) but only some rows (genes) of the original matrix. Therefore, clustering is achieved directly when instances with similar pattern of responses are grouped together.

2.3 Bi-dimensional multivariate regression trees

Now consider a multivariate response situation where each condition can also be described by a set of attributes, W_1, \dots, W_q , just as the instances can be described by attributes X_1, \dots, X_p . In this case, the response matrix $\mathcal{Y} = (y_{ij})$ can be transposed, and the regression problem can be defined for the conditions. Each condition is now treated as an instance. The observations for the j th condition can be written as $(\mathbf{w}_j, \mathbf{y}_j)$, where $\mathbf{w}_j = (w_{j1}, \dots, w_{jq})$ and $\mathbf{y}_j = (y_{1j}, \dots, y_{nj})$, with n being the number of conditions and q the number of attributes of a condition. A regression tree can then be learned to model the condition data. For clarity, we call the conditions *column instances* and the instances *row instances*, since they correspond to the columns and rows of the response matrix, respectively. Accordingly, we call X_1, \dots, X_p *row attributes*, and W_1, \dots, W_q *column attributes*.

In the case of gene expression analysis, the column attributes that can be used to describe each condition are the expression levels of a set of candidate regulators under that condition (Fig. 1, boxes C and D). Therefore, a regression tree built from the column instances explains the expression levels of genes under different conditions with the expression levels of selected regulators. The motivating assumption is that the expression level of a gene depends on the expression levels of its regulators.

The goal of our method is to model the responses using both row attributes and column attributes, i.e., to find the row attributes and column attributes that can explain the responses. In the regression tree framework, this corresponds to recursively partitioning the response matrix horizontally according to row attributes and vertically according to column attributes. The objective is to make the sub-matrices in child nodes as homogeneous as possible with respect to responses.

Formally, the input to the algorithm includes a response matrix $\mathcal{Y} = (y_{ij})$, the associated row attribute matrix $\mathcal{X} = (\mathbf{x}_i)$, and column attribute matrix $\mathcal{W} = (\mathbf{w}_j)$, where $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$ is the set of attributes for the i th row, $i = 1, \dots, m$, and $\mathbf{w}_j = (w_{1j}, \dots, w_{qj})$ is the set of attributes for the j th column, $j = 1, \dots, n$. For example, the three matrices in Fig. 1 representing motif scores, gene expression levels and regulator expression levels correspond to \mathcal{X} , \mathcal{Y} and \mathcal{W} , respectively. Each split divides the response matrix \mathcal{Y} vertically or horizontally. As a result, each node of the regression tree contains a submatrix of \mathcal{Y} and the associated row and column attributes.

To facilitate subsequent discussions, we define some symbols and notations. Let $t(r, c)$ denote a node of a tree, where r contains a subset of the row indices $\{1, \dots, m\}$, and c contains a subset of the column indices $\{1, \dots, n\}$. When there is no confusion, we abbreviate $t(r, c)$ as t or (r, c) . Let n_r and n_c denote the number of rows and the number of columns of the response matrix in node t , respectively. Let \mathbf{Y}^t denote the response matrix in t . Let $\bar{\mathbf{y}}_{r*}^t$ denote the vector of average responses across all columns in t and $\bar{\mathbf{y}}_{*c}^t$ the vector of average responses across all rows in t . Furthermore let \bar{y}_{i*}^t denote the i th element of $\bar{\mathbf{y}}_{r*}^t$ and \bar{y}_{*j}^t the j th element of $\bar{\mathbf{y}}_{*c}^t$. Let \bar{y}^t denote the average response of all rows and columns in t . The superscript t is dropped when there is no confusion.

A critical issue in our algorithm is to design a measure to reflect the homogeneity of the response matrix on both dimensions. For a node $t(r, c)$, a good candidate is the sum of squared residues introduced by Cheng and Church (2000) for bi-clustering:

$$SS2(r, c) = \sum_{i \in r, j \in c} (y_{ij} - \bar{y}_{i*} - \bar{y}_{*j} + \bar{y})^2, \quad (4)$$

where the superscript t on variables has been dropped. Equation (4) can also be rewritten as:

$$SS2(r, c) = n_r n_c (s^2(\mathbf{Y}) - s^2(\bar{\mathbf{y}}_{r*}) - s^2(\bar{\mathbf{y}}_{*c})), \quad (5)$$

where $s^2(\mathbf{Y}) = \sum_{i,j} (y_{ij} - \bar{y})^2 / n_r n_c$ is the sample variance of \mathbf{Y} , $s^2(\bar{\mathbf{y}}_{r*}) = \sum_i (\bar{y}_{i*} - \bar{y})^2 / n_r$ is the sample variance of the vector $\bar{\mathbf{y}}_{r*}$, and $s^2(\bar{\mathbf{y}}_{*c}) = \sum_j (\bar{y}_{*j} - \bar{y})^2 / n_c$ is the sample variance of the vector $\bar{\mathbf{y}}_{*c}$. It can be seen that $SS2(r, c)$ is minimal if the variance of the matrix can be explained by the variance of the row averages and the variance of the column averages. The expected value of $SS2(r, c)$ is a function of n_r and n_c :

$$E(SS2(r, c)) = (n_r n_c - n_r - n_c) \delta^2(\mathbf{Y}), \quad (6)$$

where $\delta^2(\mathbf{Y})$ is the population variance of the responses. A proof of Equation (6) by the central limit theorem is provided on the Supplementary website <http://cic.cs.wustl.edu/bdtree/>.

With $SS2(r, c)$ defined, the gain of a split can be calculated the same as in Equation (2). Note that when a row split is taken, the average responses across columns are not affected, and vice versa for column splits. Therefore, when the response matrix is split horizontally, with r_1 and r_2 rows in each child node, respectively, the gain can be computed by:

$$G2_r(r, r_1, r_2) = n_c (n_{r_1} s^2(\bar{\mathbf{y}}_{*c}^1) + n_{r_2} s^2(\bar{\mathbf{y}}_{*c}^2) - n_r s^2(\bar{\mathbf{y}}_{*c})), \quad (7)$$

where $\bar{\mathbf{y}}_{*c}^i$, $i = 1, 2$, is the $1 \times n_c$ vector of average responses across all rows in child node $t(r_i, c)$. Similarly, when the response

matrix is split vertically, with c_1 and c_2 columns in each child node, respectively, the gain can be computed by:

$$G2_c(c, c_1, c_2) = n_r(n_{c_1}s^2(\bar{\mathbf{y}}_{r*}^1) + n_{c_2}s^2(\bar{\mathbf{y}}_{r*}^2) - n_c s^2(\bar{\mathbf{y}}_{r*})), \quad (8)$$

with $\bar{\mathbf{y}}_{r*}^i$ similarly defined.

Given Equations (7) and (8) for calculating gains, the algorithm tests all possible splits on rows and columns, and select the one with the highest gain. Therefore, the algorithm automatically determines whether the split should be done on rows or columns. However, when the shape of the initial response matrix \mathcal{Y} is skewed, the gain function prefers the split on the longer side. For example, if $m \gg n$ (which is normally the case in gene expression data), the gain function prefers to split on columns to produce even narrower submatrices. This is because the expected gain for a split on one dimension is proportional to the length of the other dimension (see the Supplementary website for a proof):

$$\begin{cases} E(G2_r(r, r_1, r_2)) \simeq n_c \delta^2(\mathbf{Y}); \\ E(G2_c(c, c_1, c_2)) \simeq n_r \delta^2(\mathbf{Y}). \end{cases} \quad (9)$$

To scale down this systematic bias, we define and calculate the following adjusted gains instead:

$$\begin{cases} G2'_r(r, r_1, r_2) = G2_r(r, r_1, r_2)/n_c; \\ G2'_c(c, c_1, c_2) = G2_c(c, c_1, c_2)/n_r. \end{cases} \quad (10)$$

A problem with the homogeneity measure by *SS2* is that the produced clusters are often not tight. For example, the *SS2* measure of a matrix is zero if all rows (or columns) differ only by some constant values, i.e., $\mathbf{y}_i - \mathbf{y}_j = a_{ij}\mathbf{i}$, for all i, j , where \mathbf{y}_i and \mathbf{y}_j are the i th and j th rows of \mathbf{Y} , respectively, and $\mathbf{i} = (1, \dots, 1)$.

Consider a matrix $\begin{vmatrix} -3 & -2 & -1 \\ -1 & 0 & 1 \\ 1 & 2 & 3 \end{vmatrix}$, where each row is a gene and the values are the log ratios of expression levels under different conditions. The cluster does not seem to provide any biological significance, despite a perfect score. To deal with this problem, we define different sum-of-squares for rows and columns:

$$\begin{cases} SS3_r(r, c) = \sum_{i \in r, j \in c} (y_{ij} - \bar{y}_{*j})^2 = n_r n_c (s^2(\mathbf{Y}) - s^2(\bar{\mathbf{y}}_{*c})); \\ SS3_c(r, c) = \sum_{i \in r, j \in c} (y_{ij} - \bar{y}_{i*})^2 = n_r n_c (s^2(\mathbf{Y}) - s^2(\bar{\mathbf{y}}_{r*})). \end{cases} \quad (11)$$

The gain for a row or column split is defined correspondingly:

$$\begin{cases} G3_r(r, r_1, r_2) = SS3_r(r, c) - SS3_r(r_1, c) - SS3_r(r_2, c); \\ G3_c(c, c_1, c_2) = SS3_c(r, c) - SS3_c(r, c_1) - SS3_c(r, c_2). \end{cases} \quad (12)$$

$SS3_r$ or $SS3_c$ is equivalent to the sum of *SS* for each column or row, respectively, according to Equation (1). This measurement requires a good cluster to be coherent on at least one dimension of the matrix. Furthermore, the different gain functions defined for row and column splits enforce that a good row split must improve the homogeneity along the columns, and a good column split must improve the homogeneity along the rows. This enforcement is consistent with biological intuitions. For example, when a motif is used to separate two sets of genes, the genes within each set should have similar expression levels under the same conditions, while the expression levels under different conditions may be different.

Interestingly, the expected values of *G3* are the same as in Equation (9). Therefore, a systematic bias between G_r and G_c still exists, and the adjusted gain can be defined similarly as in Equation (10).

Given the gain functions, the algorithm proceeds the same as in the case of a single response described in Section 2.1, except the second step:

- 2'. If the current node has not met the stopping criterion, examine every possible binary split of the row instances or the column instances within the node, based on each row attribute X_i , $i = 1, \dots, p$, or column attribute W_j , $j = 1, \dots, q$, respectively, such that the attribute values for all the instances in one subset are smaller than those in the other subset.

When analyzing gene expression data, the column attributes are regulators, which may also appear in the list of genes in the rows. The algorithm does not allow a regulator to be the splitting attribute of a node that contains the regulator itself, since a gene's expression level can always be used to predict its own expression.

To prevent the tree from over-fitting the data, several parameters are implemented to control the tree size, including the minimum gain required to split a node, the minimum numbers of rows and columns within a leaf node, and the maximum number of leaf nodes. In addition, a post-pruning procedure can be performed with a separate test set, where an internal node is converted to a leaf node if by doing so the prediction accuracy on the test data does not decrease.

To predict a response, the corresponding row and column attributes are compared to the threshold values at each tree node and a branch is taken according to the result of the comparison at each step. Starting from the root node, the algorithm will always end at a terminal node t . The average value of the elements in \mathbf{Y}^t is used as the predicted value.

2.4 Cross-validation and functional analysis

The prediction accuracy of BDTree is estimated by cross-validations. The procedure of cross-validation in BDTree is slightly different from that in a one-dimensional method. Given a training data set, we denote the set of row instances as r and column instances as c . To perform a ten-fold cross-validation, r and c are both randomly divided into 10 subsets of roughly equal size, denoted by r_1, \dots, r_{10} and c_1, \dots, c_{10} , respectively. Every time a submatrix containing nine subsets of the rows and nine subsets of the columns, $(r \setminus r_i) \times (c \setminus c_i)$, is used for training, while three submatrices, $r_i \times (c \setminus c_i)$, $(r \setminus r_i) \times c_i$, and $r_i \times c_i$ are used for testing, for $i = 1, \dots, 10$. The mean squared errors or the correlation coefficient between the predicted and actual values are calculated as a measure of accuracy. In addition, accuracies can be calculated for the three testing submatrices separately, corresponding to the prediction accuracy for unseen rows, unseen columns, and unseen rows plus unseen columns.

In the case of analyzing gene expression data, each leaf node of the tree contains a subset of the genes and a subset of the conditions. To determine the functional relevance of the splits, we calculate the enrichment of gene ontology (GO) terms (Harris *et al.*, 2004) within each leaf node. When possible, we also group the experimental conditions into categories, and calculate the enrichment of particular categories within each leaf node. The significance of enrichment is measured by an accumulative hyper-geometric test, and the p -values are adjusted by Bonferroni corrections for multiple tests (Altman, 1991).

When a tree is built, it automatically selects a set of attributes from the row attributes and column attributes to explain the pattern of responses. Intuitively, the gain of splitting a node with a certain attribute can be used as a measure of the importance of the attribute. However, the tree only selects the attribute with the highest gain at each step, while ignoring all the others. Phuong *et al.* (2004) provides a better method to measure the importance of all attributes, based on surrogate splits. We adopt the same idea, but rank row and column attributes separately.

3 RESULTS

3.1 Data sets

To demonstrate the strength of our method in identifying real motifs and regulators, we tested it on two microarray sets for *S. cerevisiae*. The cell-cycle data set consists of the expression levels of 800 genes measured under 77 different time points in cell cycles from several experiments (Spellman *et al.*, 1998). The stress response data set includes gene expression data collected under 173 different stress conditions. As in Middendorf *et al.* (2004), we selected 1411 genes from the stress data set, which includes 469 highly variant genes and 1250 genes that are part of the 17 clusters identified by Gasch *et al.* (2000). We downloaded the background normalized log ratios of all genes from SGD (Dwight *et al.*, 2004).

We used the set of 466 candidate regulators compiled by Segal *et al.* (2003) as column attributes. These include both TFs and signaling molecules that may have transcriptional impact. We combined three types of row attributes: a set of 356 motifs from Pilpel *et al.* (2001), binding data of 204 TFs measured under various environmental conditions (Harbison *et al.*, 2004), and 615 over-represented k -mers ($5 \leq k \leq 7$) discovered by a steganalysis-based motif finding algorithm called WordSpy developed in our lab (Wang *et al.*, 2005). The set of motifs from Pilpel *et al.* includes known motifs and putative motifs identified computationally from common functional categories. We used RSA tools (van Helden, 2003) to retrieve up to 500bp of intergenic sequences upstream of each gene start codon as its promoter, and searched both strands for the appearance of motifs and k -mers. The program ScanACE (Roth *et al.*, 1998) was used to scan each promoter, and the highest score for each motif was recorded as its attribute value. For each k -mer, its number of occurrences in a promoter sequence is used as its score. All data including gene expression and attributes are available on the Supplementary website.

3.2 Simulation study of gain functions

We have shown in Section 2.3 that for both SS2 and SS3, the theoretically expected gain resulted from a random split is correlated with the size of the dimension that is unsplit. Here, we use simulation to show that this is also true in practice.

We first considered the case where gene expression levels are identically and independently distributed (i.i.d.). We randomly shuffled the yeast cell-cycle gene expression matrix, which is then split into two sub-matrices by randomly dividing its rows or columns into two sets. The relative sizes of the two sub-matrices vary from 1:9 to 5:5. Fig. 2(a) shows the average gains of 1000 random splits on rows or columns, calculated using SS2. As shown, the average gains agree with the theoretical results almost perfectly, regardless the relative sizes of the two sub-matrices. The gains of

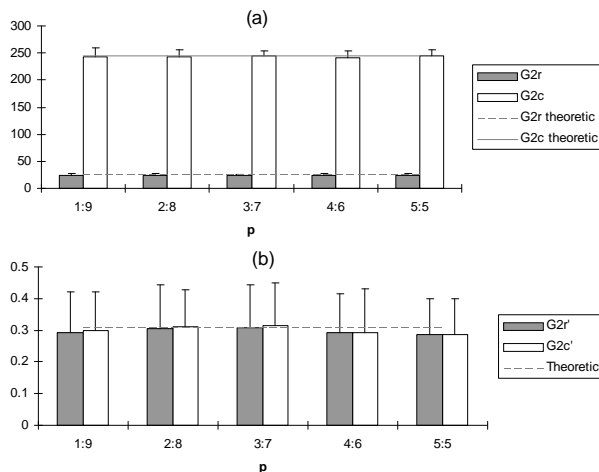


Fig. 2. Simulated gains. (a) Unadjusted gains on an i.i.d. matrix; (b) Adjusted gains on a real matrix. P is the relative sizes of the two sub-matrices after splitting.

column splits are much larger than those of row splits, which justifies the adjustment of gains by Equation (10). Next, we repeated the experiments on the real yeast cell-cycle expression matrix, to which the i.i.d. assumption does not hold. The average gains are close to the theoretical values, although the agreement is not as well as in the i.i.d. case. As shown in Fig. 2(b), the adjusted gains resulted from row splits or column splits have similar means and standard deviations. The results using SS3 or the stress response data set are similar.

3.3 Model accuracy

To evaluate the performance of our method, we applied it to the yeast stress response dataset, and conducted ten-fold cross-validations. We calculated the correlation coefficients between the predicted and actual values as a measure of accuracy.

In the first set of experiments, we analyzed the effect of the choices of parameters. We have shown by simulation that the adjustment of gains is necessary to eliminate the systematic bias between row and column splits. In fact, the model built with adjusted gains has a higher cross-validation accuracy than that with unadjusted gains (0.54 vs. 0.43). We also found that the tree based on SS3 has better accuracy than that based on SS2 (0.54 vs. 0.45). In addition, the prediction accuracy for unseen genes (0.56) is slightly higher than for unseen conditions (0.52) or unseen genes plus unseen conditions (0.51).

Next, We compared the accuracy of BDTree to the k -nearest neighbor (KNN) method. With the KNN method, the expression level of a gene at a certain condition was predicted by the average expression level of the k nearest genes under the k nearest conditions (best $k = 20$ in our experiment), where the distance between genes or conditions was defined by the Euclidean distance of their normalized attribute vectors. We chose KNN as a base level classifier because it is relatively easy to implement a bi-dimensional counterpart of our algorithm. The cross-validation accuracy of BDTree (0.54) is much higher than that of the KNN method (0.37).

Table 1. confusion matrix

		Predicted by BDTree			Predicted by GeneClass		
		Down	Baseline	Up	Down	Baseline	Up
TRUE	Down	15.2%	7.1%	2.9%	16.5%	8.9%	1.5%
	Baseline	7.0%	34.5%	9.2%	9.3%	32.4%	6.3%
	Up	2.4%	8.4%	13.2%	2.8%	9.9%	12.0%

Second, we considered the case where BDTree is grown using row attributes only. This is equivalent to the method of Phuong *et al.* (2004) in spirit. Since the expression matrix is only partitioned horizontally, the method is unable to predict expression levels under unseen conditions. Therefore, we conducted cross-validations only on unseen genes. The correlation coefficient obtained by this method is similar to our full model where both row and column attributes are used (0.57 vs. 0.56), which means that our method did not lose any information in row attributes even though column attributes are used together.

Finally, we compared our method to the GeneClass method (Middendorf *et al.*, 2004), which is similar to ours in that it can also predict gene expression levels using both regulators and binding motifs. One problem when comparing to their results, however, is that their method is only applicable to pre-discretized expression levels, while our method can be applied to real expression values. Therefore, we discretized the expression data into three levels (up, down and baseline) as in their method, and obtained a confusion matrix for our predictions as shown in Table 1. It turned out that the two methods have similar prediction accuracies (63% for ours vs. 61% for theirs) using discretization. On the other hand, their method has used a technique called boosting, which greatly improves prediction accuracy, but reduces interpretability of models. Besides not requiring discretization, our method also has several additional advantages that will be discussed in Section 4.

3.4 Biological interpretation and functional analysis

Fig. 3 shows the regression tree learned from the yeast stress response data and the gene expression matrix reorganized according to the tree. The interpretation of the tree is straightforward. Each oval represents a row attribute (prefixed by "h.") for horizontal splitting, or column attribute (prefixed by "v.") for vertical splitting. The edge labels represent the thresholds used for splitting. Each gray box is a leaf node, where the first label is the ID of the node, and the second label is the average gene expression for the sub-matrix corresponding to the node. Note that two subtrees are collapsed to save space (shown as N373 and N785). Each path from the root node to a leaf node forms a rule, which represents a biological hypothesis about the logic relationships among the expression levels of target genes, presence of binding motifs, and the expression levels of putative regulators. For example, node N1267 shows that the expression of genes with both mRRPE and PAC motifs are highly repressed if the expression of Tpk1 and Kin82 are both induced.

In order to better interpret and understand these rules, we calculated the enrichment of GO functional categories for the genes within each leaf node. We also grouped the 173 experimental conditions into 19 categories, and identified the significantly enriched categories for each leaf node. Together, this describes the functional

roles of a set of genes, their regulators and motifs, and the conditions under which they are activated or deactivated. The complete results of the analysis can be viewed interactively on the Supplementary website by clicking on the tree nodes. Overall, among the 50 leaf nodes, 45 have enriched conditions and 42 have enriched GO categories with corrected p -values < 0.05 (see Section 2.4).

On the top levels, the matrix is horizontally partitioned into five blocks by row attributes (binding motifs or binding factors). The first block, corresponding to leaf nodes from N1084 to N1267, contains 245 genes whose promoters all have mRRPE motifs. It is known that mRRPE is important in regulating rRNA transcription and processing, a process that is repressed under many stress conditions (Pilpel *et al.*, 2001; Gasch *et al.*, 2000). Indeed, GO analysis showed that 74 of these genes participate in rRNA processing ($p < 8e-65$). Furthermore, 94 genes having both mRRPE and PAC motifs (nodes N1252 and N1267) are more drastically repressed than the genes having mRRPE only, which agrees with the fact that the two motifs work cooperatively (Sudarsanam *et al.*, 2002). Comparing the GO annotation of the genes having both motifs and the genes having mRRPE only, the former are much more enriched in nucleolus (63/94 vs. 32/151, $p < 8e-13$), and ribosome biogenesis and assembly (60/94 vs. 36/151, $p < 5e-10$). On the other hand, the genes having only mRRPE motifs are more enriched in cytoplasm (95/151 vs. 22/94, $p < 1e-9$), and protein biosynthesis (38/151 vs. 6/94, $p < 9e-5$). This suggests that the genes regulated by both mRRPE and PAC have regulatory roles in ribosome biosynthesis, while the genes regulated by mRRPE alone are involved in protein synthesis. Interestingly, the genes having two or more copies of ttctt are down-regulated (N1219), while the genes having two or more copies of tttct are up-regulated (N1214). Unlike other leaves in this block, N1219 is enriched in polysaccharide metabolism/glycan metabolism ($p < 0.003$). This suggests that, although ttctt and tttct are similar to each other, they may be binding motifs of different TFs (Cliften *et al.*, 2003).

The other four blocks also provide some biological insights. The second block contains 105 genes regulated by Fhl1 (leaf nodes from N1026 to N1064). Among them, 97 are structural constituent of ribosome ($p < 3e-144$). This is consistent with the recent results that Fhl1 (together with Ifh1) plays a central role in ribosome protein gene regulation (Wade *et al.*, 2004). The third block contains leaf nodes from N579 to N1015. The main binding motif in this block, agggg, is the consensus sequence of the stress response element (STRE), which induces a large number of stress-responsive genes (Gasch *et al.*, 2000). The most enriched GO category is generation of precursor metabolites and energy (60/401, $p < 3e-26$). The genes having three or more copies of STRE motifs (node N986) show higher inductivity than those having less copies, and are enriched in energy reserve metabolism ($p < 1e-9$), or specifically, trehalose metabolism, which is an important determinant of stress resistance in yeast (Winderickx *et al.*, 1996). This suggests that the TF bound to STRE may have a preference for repetitive motifs. The fourth block (N526, N539 and N550) has 76 genes that are regulated by ctactg, which is the consensus sequence of PAC. Like the genes in N1267, this set of genes are also enriched in ribosome biogenesis and assembly ($p < 1e-10$), although the degree of enrichment lower than in N1267 (20/76 vs. 60/94, $p < 1e-6$). Lastly, the fifth block is enriched with genes for Nitrogen compound metabolism ($p < 7e-21$). Note that to the right of this block, the genes are separated into many small subsets, each of which is regulated by

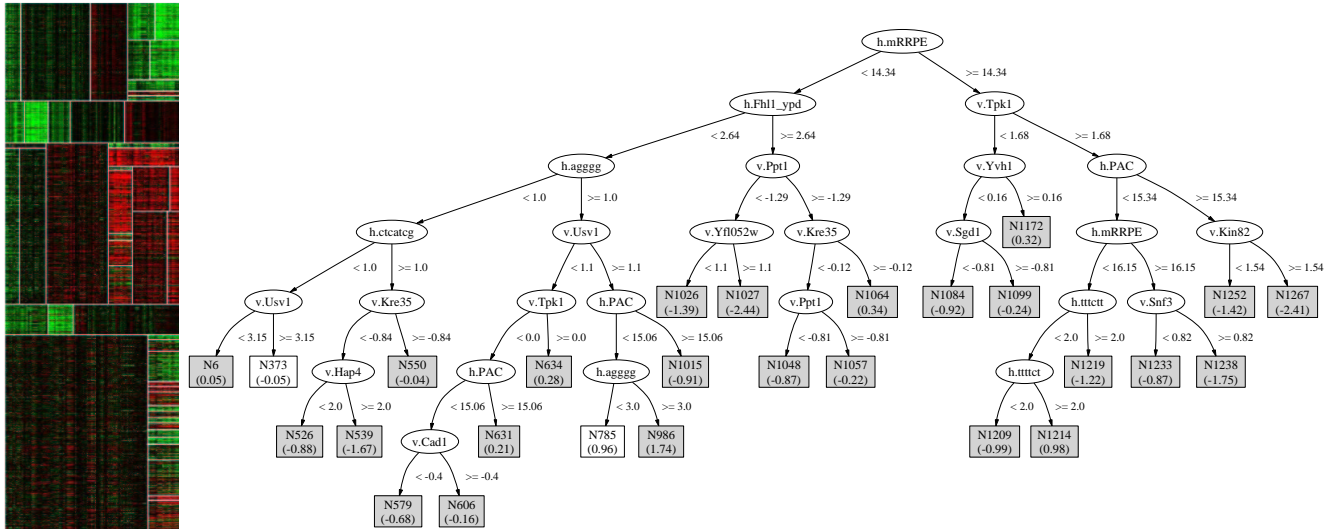


Fig. 3. Model built by BDTree for the yeast stress-responsive genes and the partition of the expression matrix. Node labels prefixed by “h.” (“v.”) are row (column) attributes. Row attributes whose names start with a upper-case letter followed by lower-case letters are from binding data. Row attributes with all lower-case letters names are *k*-mers from wordspy. The remaining row attributes are from the motif set of Pilpel *et al.* The bottom-left and upper-right submatrices correspond to leaf nodes N6 and 1267, respectively. This figure can be viewed with a higher resolution on the Supplementary website.

a different motif. GO analysis revealed that each subset of genes is responsible for the metabolism of different substrates. For example, there are nodes enriched with asparagine catabolism (N392 on the Supplementary website, $p < 3e-8$), aldehyde metabolism (N375, $p < 0.015$), methionine metabolism (N442, $p < 0.002$) and glucosidase activity (N521, $p < 0.001$).

On the other hand, although we found that most regulators selected by the tree have been reported as important in regulating stress-responsive genes (e.g., Usv1, Ppt1, Tpk1, Kin82), the exact biological role of putative regulators is hard to determine. One reason is that TFs are often post-transcriptional modified or translocated before it can bind to a promoter. Therefore, the mRNA levels of a TF may not indicate its activity. In some cases, since we have included signaling molecules as candidate regulators, our method may select a kinase that activates the TF instead of the TF itself. In general, when the transcription of a gene requires a cascade of activation/deactivation, our method may select a regulator upstream to the actual transcriptional regulation. Nevertheless, we found literature support for some of the rules. For example, the tree shows that the expression of Tpk1, a subunit of cAMP-dependent protein kinase (PKA), is negatively associated with the expression of mRRPE targets. It has been reported that cAMP signaling pathway regulates the down-regulation of ribosome biogenesis (Schawalter *et al.*, 2004). Furthermore, it is known that the RAS/cAMP pathway negatively regulates cellular physiology characteristic of stationary-phase (Schawalter *et al.*, 2004). This agrees with our results that node N1267 is enriched with GO annotation ribosome biogenesis and assembly, and the most significant conditions are “stationary phase”. Another computational study by Segal *et al.* (2003) suggested a role for Tpk1 and Sgd1 in the regulatory program for rRNA processing and ribosome biogenesis, which is similar to our results. Msn2, the TF that binds to STRE, needs to be translocated from the cytoplasm to nucleus under stress conditions, and it has been reported that PKA is involved in the translocation (Jacquet *et al.*, 2003),

which is consistent with our analysis (node N634). Our results suggest that Usv1, which has been identified as a top regulator for many stress response (Segal *et al.*, 2003), may play a role in this process as well.

We also calculated the importance measure of each attribute using surrogate splits (see Section 2.4). Table 2 shows the top 20 row attributes and top 20 column attributes respectively. The complete list is available on the Supplementary website.

Among the top 20 row attributes, PAC (as well as gctcatc and ctcatcg), mRRPE, and Fhl1 are known to be related to stress response as discussed above. agggg, ccctt, ggggc and aggggc are variants of the extended STRE motif (a/caggggc/ggg) or its reverse complement (Harbison *et al.*, 2004). tcctt and tcctt are the binding motifs of Gis1, a transcriptional factor involved in the expression of genes during nutrient limitation (Oshiro *et al.*, 2003). Rap1 and Snf1 are known to control the expression of ribosomal protein genes during various stress responses (Gasch *et al.*, 2000; Dwight *et al.*, 2004). Both Gat3 and Yap5 have functions in stress responses and co-bind with Msn4 (Banerjee and Zhang, 2003). The binding data of Fhl1 and Rap1 measured under different conditions are all top-ranked row-attributes, which means that their binding is probably condition invariant. On the other hand, STRE-like motifs rather than Msn2 are ranked as top row attributes, which suggests that the binding specificity of Msn2 varies significantly under different conditions.

Many of the top 20 column attributes, such as Usv1, Tpk1, Xbp1, Gis1, Kin82, Gac1, Rim11, Gpa2, Yj1103c, and Tos8, have evidence to be involved in regulating various stress responses in SGD database (Dwight *et al.*, 2004) or other computational analyses (Middendorf *et al.*, 2004; Segal *et al.*, 2003). Interestingly, we find that only a few of the identified top regulators are TFs (XBPI, GIS1, TOS8, NRG1, GAT2), while the majority are protein kinases or hydrolases. This suggests that the activities of many stress-responsive TFs are regulated post-transcriptionally.

Table 2. Top row and column attributes identified by BDTree

top row attributes	PAC, mRRPE, Fhl1_rapa, Fhl1_ypd, agggg, Fhl1_sM, Fhl1_h2o2hi, ccctt, Rap1_sm, gctcatc, Rap1_ypd, Gat3_ypd, Rap1, ggggc, Sfp1_sm, ctcatcg, aggggc, tcctt, Yap5_ypd, tcctt
top column attributes	Usv1, Tpk1, Xbp1, Kns1, Sip2, Kin82, Yjl103c, Mtl1, Ppz2, Yak1, Gis1, Pde1, Rim11, Gpa2, Tpk2, Tos8, Nrg1, Gat2

Row attributes starting with upper-case letter followed by lower-case letters are from binding data of Harbison *et al.*. Row attributes with all upper-case letters are from the motif set of Pilpel *et al.*. rapa: nutrient deprived. sm: amino acid starvation. ypd: normal growth condition.

We also learned a model of the yeast cell-cycle data. The tree and detailed analysis is on the Supplementary website. Our method identified almost all known TFs regulating the yeast cell-cycle genes and their binding motifs (see Supplementary website).

4 DISCUSSION

In this research, we have developed a novel method, the bi-dimensional regression tree (BDTree), for modeling transcriptional regulation from large-scale gene expression data. BDTree is a significant extension of previous works. First, the tree-based approach does not assume linear additivity of regulatory elements or any distribution of the underlying data set. Second, by considering gene expression under multiple conditions simultaneously, the method can tolerate more noises than using individual arrays. More importantly, by taking into account both the expression of putative regulators and the occurrence of putative binding motifs, BDTree is able to identify *condition-specific* regulatory elements and regulators for each gene. We have successfully applied BDTree to the yeast cell-cycle and stress response data, and identified many biologically significant binding motifs and regulators.

Two existing methods are similar to our approach in that they also attempt to model the large-scale gene expression data under a large number of conditions. The module networks approach (Segal *et al.*, 2003) clusters genes according to their expression patterns, and builds a regression tree for each cluster. However, in their method, binding motifs are not considered when clustering genes. As a result, genes having similar expression patterns are assigned the same set of regulators, regardless of the difference of their promoters. Furthermore, the clustering of genes in their method is based on the expression levels across all conditions. Therefore, their method is unlikely to identify condition-specific regulation.

The GeneClass method (Middendorf *et al.*, 2004) is the most similar to ours. Indeed, GeneClass and BDTree have the same schematic representation (Fig. 1 box E), i.e., modeling gene expression levels from putative binding motifs and TF expression levels. In addition, both methods build tree-based models (decision trees in GeneClass and regression trees in BDTree). Despite these similarities, the underlying modeling rationales are very different. BDTree is a novel extension to a multivariate regression tree approach, while GeneClass transforms the modeling problem into a traditional decision tree learning problem. This difference leads to several significant consequences.

First, in GeneClass, gene expression levels have to be discretized into three categories: up, down, and intermediate, but only the up

and down categories are used for training. These choices are arbitrary and may cause a significant amount of information to be lost. BDTree, on the other hand, accepts real-valued data and uses all data points. Second, because GeneClass treats expression levels as univariate variables, the differences between genes and conditions are disregarded. As a result, GeneClass attempts to find submatrices that have constant expression levels along both dimensions, which may not be biologically meaningful. In contrast, BDTree optimizes the homogeneity on one dimension of the expression matrix in each split, while allows heterogeneity on the other dimension. Third, GeneClass forces each split to couple a row attribute and a column attribute. Therefore, a total of mn attribute pairs need to be considered for each split, where m and n are the numbers of row and column attributes, respectively. BDTree only needs to consider $m + n$ row and column attributes individually for each split and is thus more scalable. By coupling row and column attributes, GeneClass may have the advantage of directly suggesting associations between regulators and binding motifs. However, it is not always advisable to relate regulators to binding motifs. For example, the binding motifs of a regulator or a regulator itself may not be present in the list of candidates. It is also possible that the regulators are post-transcriptionally regulated; therefore its expression levels do not correlate with the expression levels of its targets. BDTree is more flexible since it does not force an explicit pairing of regulators and binding motifs.

Our method is general and can be turned into several previous methods easily, taking each of them as a special case. For example, when the minimum number of column instances is set to a sufficiently large number, BDTree is equivalent to that in Phuong *et al.* (2004). When vertical splits are restricted to occur only after horizontal splits are completed, BDTree performs similarly as the method of Segal *et al.* (2003). BDTree can also be applied to domains other than computational biology, such as clinical studies. For example, the multivariate responses may be a time-series observation of drug efficacies on patients, for which our method can be used to identify the time-dependent impact of different factors.

There are several directions that this method may be extended. One problem with the current implementation of BDTree is that the splitting of genes or conditions is strictly based on attribute values. Such a hard split may be undesirable considering that there are inevitable noises in attribute values: motif representations may be inaccurate, expression levels of regulators are unreliable, and normalization may introduce additional noises. Furthermore, the regression tree learning algorithm is essentially greedy, never re-considering a split that has been made. To circumvent these problems, some fuzzy rules may be applied to find soft splitting. Look-ahead strategies may be used to find globally better top-level splits. Furthermore, some iterative strategies may be adopted to improve motif representations. Another problem with our approach is that it can only model simultaneous changes between regulators and targets, i.e., the expression levels of regulators have to be correlated or anti-correlated with those of the target genes. To allow shifted or reversely shifted correlations to be identified, the method may include the expression levels of regulators at previous time points as additional column attributes.

ACKNOWLEDGEMENT

This research was supported in part by NSF grants IIS-0196057 and EIA-0113618 under the ITR program, and a grant from Monsanto Corporation.

REFERENCES

- Altman, D. (1991) *Practical Statistics for Medical Research*. Chapman & Hall/CRC.
- Banerjee, N. and Zhang, M. (2003) Identifying cooperativity among transcription factors controlling the cell cycle in yeast. *Nucleic Acids Res.*, **31**, 7024–31.
- Beer, M. and Tavazoie, S. (2004) Predicting gene expression from sequence. *Cell*, **117** (2), 185–98.
- Breiman, L., Friedman, J. and Stone, R.O.C. (1984) *Classification and Regression Trees*. Wadsworth Int. Group, Belmont, CA.
- Bussemaker, H., Li, H. and Siggia, E. (2001) Regulatory element detection using correlation with expression. *Nat. Genet.*, **27**, 167–71.
- Cheng, Y. and Church, G. (2000) Biclustering of expression data. *Proc Int Conf Intell Syst Mol Biol.*, **8**, 93–103.
- Cliften, P., Sudarsanam, P., Desikan, A., Fulton, L., Fulton, B., Majors, J., Waterston, R., Cohen, B. and Johnston, M. (2003) Finding functional features in saccharomyces genomes by phylogenetic footprinting. *Science*, **301** (5629), 71–6.
- Conlon, E., Liu, X., Lieb, J. and Liu, J. (2003) Integrating regulatory motif discovery and genome-wide expression analysis. *Proc. Natl. Acad. Sci. USA*, **100**, 3339–44.
- Dwight, S., Balakrishnan, R., Christie, K., Costanzo, M., Dolinski, K., Engel, S., Feierbach, B., Fisk, D., Hirschman, J., Hong, E., Issel-Tarver, L., Nash, R., Sethuraman, A., Starr, B., Theesfeld, C., Andrada, R., Binkley, G., Dong, Q., Lane, C., Schroeder, M., Weng, S., Botstein, D. and Cherry, J. (2004) Saccharomyces genome database: underlying principles and organisation. *Brief Bioinform.*, **5** (1), 9–22.
- Gasch, A., Spellman, P., Kao, C., Carmel-Harel, O., Eisen, M., Storz, G., Botstein, D. and Brown, P. (2000) Genomic expression programs in the response of yeast cells to environmental changes. *Mol Biol Cell*, **11** (12), 4241–57.
- Harbison, C., Gordon, D., Lee, T., Rinaldi, N., Macisaac, K., Danford, T., Hannett, N., Tagne, J., Reynolds, D., Yoo, J., Jennings, E., Zeitlinger, J., Pokholok, D., Kellis, M., Rolfe, P., Takusagawa, K., Lander, E., Gifford, D., Fraenkel, E. and Young, R. (2004) Transcriptional regulatory code of a eukaryotic genome. *Nature*, **431**, 99–104.
- Harris, M., Clark, J., Ireland, A., Lomax, J., Ashburner, M., Foulger, R., Eilbeck, K., Lewis, S., Marshall, B., Mungall, C., Richter, J., Rubin, G., Blake, J., Bult, C., Dolan, M., Drabkin, H., Eppig, J., Hill, D., Ni, L., Ringwald, M., Balakrishnan, R., Cherry, J., Christie, K., Costanzo, M., Dwight, S., Engel, S., Fisk, D., Hirschman, J., Hong, E., Nash, R., Sethuraman, A., Theesfeld, C., Botstein, D., Dolinski, K., Feierbach, B., Berardini, T., Mundodi, S., Rhee, S., Apweiler, R., Barrell, D., Camon, E., Dimmer, E., Lee, V., Chisholm, R., Gaudet, P., Kibbe, W., Kishore, R., Schwarz, E., Sternberg, P., Gwinn, M., Hannick, L., Wortman, J., Berriman, M., Wood, V., de la Cruz, N., Tonellato, P., Jaiswal, P., Seigfried, T. and White, R. (2004) The gene ontology (GO) database and informatics resource. *Nucleic Acids Res.*, **32** (Database issue), D258–61.
- Hu, Y., Sandmeyer, S., McLaughlin, C. and Kibler, D. (2000) Combinatorial motif analysis and hypothesis generation on a genomic scale. *Bioinformatics*, **16** (3), 222–32.
- Jacquet, M., Renault, G., Lallet, S., Mey, J.D. and Goldbeter, A. (2003) Oscillatory nucleocytoplasmic shuttling of the general stress response transcriptional activators Msn2 and Msn4 in *Saccharomyces cerevisiae*. *J Cell Biol.*, **161** (3), 497–505.
- Keles, S., van der Laan, M. and Eisen, M. (2002) Identification of regulatory elements using a feature selection method. *Bioinformatics*, **18**, 1167–75.
- Middendorf, M., Kundaje, A., Wiggins, C., Freund, Y. and Leslie, C. (2004) Predicting genetic regulatory response using classification. *Bioinformatics*, **20** (Suppl 1), I232–I240.
- Oshiro, J., Han, G., Iwanyszyn, W., Conover, K. and Carman, G. (2003) Regulation of the yeast DPP1-encoded diacylglycerol pyrophosphate phosphatase by transcription factor Gis1p. *J Biol Chem.*, **278** (34), 31495–503.
- Phuong, T., Lee, D. and Lee, K. (2004) Regression trees for regulatory element identification. *Bioinformatics*, **20** (5), 750–7.
- Pilpel, Y., Sudarsanam, P. and Church, G. (2001) Identifying regulatory networks by combinatorial analysis of promoter elements. *Nat. Genet.*, **29**, 153–9.
- Roth, F., Hughes, J., Estep, P. and Church, G. (1998) Finding DNA regulatory motifs within unaligned noncoding sequences clustered by whole-genome mRNA quantitation. *Nat Biotechnol.*, **16**, 939–45.
- Ruan, J. and Zhang, W. (2004). Discovering transcriptional regulatory rules from gene expression and TF-DNA binding data by decision tree learning. Technical Report 43 Department of Computer Science and Engineering, Washington University in St Louis.
- Schawaldner, S., Kabani, M., Howald, I., Choudhury, U., Werner, M. and Shore, D. (2004) Growth-regulated recruitment of the essential yeast ribosomal protein gene activator Ifh1. *Nature*, **432** (7020), 1058–61.
- Segal, E., Shapira, M., Regev, A., Pe'er, D., Botstein, D., Koller, D. and Friedman, N. (2003) Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nat. Genet.*, **34**, 166–76.
- Segal, M. (1992) Tree-structured methods for longitudinal data. *J. Am. Stat. Assoc.*, **87**, 407–418.
- Soinov, L., Krestyaninova, M. and Brazma, A. (2003) Towards reconstruction of gene networks from expression data by supervised learning. *Genome Biol.*, **4**, R6.
- Spellman, P., Sherlock, G., Zhang, M., Iyer, V., Anders, K., Eisen, M., Brown, P., Botstein, D. and Futcher, B. (1998) Comprehensive identification of cell cycle-regulated genes of the yeast *Saccharomyces cerevisiae* by microarray hybridization. *Mol. Biol. Cell*, **9**, 3273–97.
- Sudarsanam, P., Pilpel, Y. and Church, G. (2002) Genome-wide co-occurrence of promoter elements reveals a cis-regulatory cassette of rRNA transcription motifs in *Saccharomyces cerevisiae*. *Genome Res.*, **12** (11), 1723–31.
- van Helden, J. (2003) Regulatory sequence analysis tools. *Nucleic Acids Res.*, **31** (13), 3593–6.
- Wade, J., Hall, D. and Struhl, K. (2004) The transcription factor Ifh1 is a key regulator of yeast ribosomal protein genes. *Nature*, **432** (7020), 1054–8.
- Wang, G., Yu, T. and Zhang, W. (2005) Wordspy: identifying transcription factor binding motifs by building a dictionary and learning a grammar. *Nucleic Acids Res.*, **33** (Web Server issue), W412–6.
- Winderickx, J., de Winde, J., Crauwels, M., Hino, A., Hohmann, S., Dijck, P.V. and Thevelein, J. (1996) Regulation of genes encoding subunits of the trehalose synthase complex in *Saccharomyces cerevisiae*: novel variations of STRE-mediated transcription control? *Mol Gen Genet.*, **252** (4), 470–82.